

ABSTRACT

Fault management and recovery in a distributed network accounting system includes classifying nodes in the system according to complexity of processing in the node. For nodes of relatively low processing complexity, context check-pointing a state of processing in the nodes to permit automatic recovery of the node to the nodes' most recent processing context checkpoint; and for nodes of relatively high complexity, producing a directed graph of the programmable nodes that controls a flow of data and control of processing through the system, and producing a dynamic modification of the directed graph to automatically fail-over to redundant back-up nodes based on thresholds established for the component hosts.

20371103.doc